Agent Sensing with Stateful Resources (Extended Abstract)

Adam Eck and Leen-Kiat Soh
Department of Computer Science and Engineering
University of Nebraska-Lincoln
256 Avery Hall Lincoln, NE 68588, USA
+1-402-472-4257
{aeck, lksoh}@cse.unl.edu

ABSTRACT

In the application of multi-agent systems to real-world problems, agents often suffer from bounded rationality where agent reasoning is limited by 1) a lack of knowledge about choices, and 2) a lack of resources required for reasoning. To overcome the former, the agent uses sensing to refine its knowledge. However, sensing can also require limited resources, leading to inaccurate environment modeling and poor decision making. In this paper, we consider a novel and difficult class of this problem where agents must use *stateful* resources during sensing, which we define as resources whose state-dependent behavior changes over time based on usage. Specifically, such sensing changes the state of a resource, and thus its behavior, producing a phenomenon where the sensing activity can and will distort its own outcome. We term this the *Observer Effect* after the similar phenomenon in the physical sciences. Given this effect, the agent faces a strategic tradeoff between satisfying the need for 1) knowledge refinement, and 2) avoiding corruption of knowledge due to distorted sensing outcomes. To address this tradeoff, we use active perception to select sensing activities and model activity selection as a Markov decision process (MDP) solved through reinforcement learning where an agent optimizes knowledge refinement while considering the state of the resource used during sensing.

Categories and Subject Descriptors

1.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence – *intelligent agents, multiagent systems*

General Terms

Performance

Keywords

Observer Effect, Bounded rationality, Stateful resources, Sensing

1. INTRODUCTION

One common problem in real-world applications of multiagent systems is **bounded rationality**. Grounded in the economics (e.g., [1,7]) and cognitive psychology (e.g., [3]) literature, this problem addresses limitations on agent reasoning. In contrast to *perfect rationality*, bounded rationality assumes agents generally lack at least one of: 1) perfect knowledge of available choices, 2) perfect knowledge of preferences over choices/outcomes, and 3) unlimited ability and resources to calculate the optimal choice.

To overcome the first two limitations, agents perform sensing to gather information about the environment and inform their decision making. However, sensing can also require limited resources, and thus sensing performance can also be affected by the agent's bounded rationality. In this paper, we consider the impact on agent sensing of one category of resources used during sensing: stateful resources. We introduce a novel side-effect from the use of this type of resource called the Observer Effect and describe our methodology for choosing sensing activities to overcome its negative consequences.

2. OBSERVER EFFECT

One important property of resources used during sensing is whether the resource is *stateless* or *stateful*. Specifically, these two types differ in the importance of resource usage history on its behavior. On the one hand, the behavior of stateless resources does not depend on the past history of their usage. For example, computational resources such as CPU cycles always process the same amount of sensed data in the same way. Stateful resources, on the other hand, behave differently depending on their past usage. For example, in a user preference elicitation scenario, a human user's patience is used up through repeated interruptions, leading to increased frustration which affects the user's cognitive workload and feelings towards the system [5]. Likewise, as an agent depletes network bandwidth, the network becomes more congested and its behavior more variable [2].

This notion of resource state is important in agent sensing because sensing actions change the underlying state of a resource, and thus, its behavior. If the outcome of the sensing activity relies on the behavior of the resource used during sensing, a phenomenon occurs where *the act of making an observation distorts the observation itself*. We term this phenomenon the **Observer Effect** (OE) after a similar phenomenon in the physical sciences. For example, in the aforementioned networking scenario, sending additional traffic to measure the network's performance reduces bandwidth which increases congestion and latency [2]. As a result, observations produced do not reflect the state of the network when sensing is not performed, thus reducing the *accuracy of information gathered by sensing*. Furthermore, in our user preference elicitation example, prompting the user with questions is an interruption which affects the user's feelings towards the system [5] which can lead to less willingness to provide responses, thus reducing the *quantity of information gathered by sensing*.

Therefore, the Observer Effect is an important problem during stateful resource-based sensing because it creates a tradeoff we call the **Observer Effect Tradeoff** between satisfying the need for 1) providing knowledge refinement to better guide its reason-

Cite as: Agent Sensing with Stateful Resources, Adam Eck and Leen-Kiat Soh. (Extended Abstract), *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1283 - 1284. Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

ing, and 2) avoiding knowledge corruption due to distorted sensing outcomes. Thus, the Observer Effect places stress on an agent's sensing activity selection for gathering information used to refine the agent's knowledge to properly achieve its goals.

Comparing resource usage during sensing and computation, we note that the state-dependent behavior of resources changed with their use during sensing results in nonmonotonic performance of sensing with respect to resource use due to the Observer Effect. In other words, while additional sensing activities which require more resource usage might lead to better knowledge refinement in some situations, this might not occur after an undesired resource state change. Thus, traditional metareasoning solutions to limited resource problems in bounded rationality such as anytime algorithms which require monotonicity [11] cannot be applied when making decisions about stateful resource use during sensing (although they have been used for stateless computational resources [10] which satisfy monotonicity). Instead, we require a solution that can handle non-monotonicity, such as the Markov decision process (MDP)-based approaches to metareasoning (e.g., [6]).

3. METHODOLOGY

Specifically, we utilize a domain-independent active perception approach to sensing [9]. From this perspective, agents actively guide sensing in order to select what information to gather, as well as how to gather and process such information, rather than passively react to whatever information the agent's sensors perceive during its task-oriented actions. To choose actions, we assume that the behavior of stateful resources is stochastic, a common assumption about the environment in multiagent systems. We also assume that the behavior of the resource depends only on its current state. Thus, we model sensing activity selection as an MDP (e.g., [4]) which we term the **Observer Effect MDP**.

In this model, the agent considers a set of *sensing states* S upon which the agent makes decisions, a set of *active perception choices* (i.e., sensing activities) A the agent can make about sensing, a function $T(s, a)$ describing the stochastic changes in sensing state based on resource usage during sensing activities, and the *amount of knowledge refinement* $R(s, a)$ produced by a sensing activity depending on the current state. Here, each sensing state is the combination of two factors impacting knowledge refinement: resource state (through the OE) and knowledge state (capacity for improvement). Using this model, the agent aims to maximize knowledge refinement $R(s, a)$ in order to handle the OE Tradeoff—by selecting sensing actions to provide positive refinement improving its knowledge while avoiding negative refinement from knowledge corruption based on the OE. Specifically, solving the Observer Effect MDP provides a policy optimizing knowledge refinement based on sensing states for sensing action selection.

Since an explicit, parameterized Observer Effect MDP model of the active perception decision process is difficult to provide *a priori* (e.g., due to environment dynamics or lack of background knowledge), we use reinforcement learning (RL) [8] to learn how to solve the OE MDP. One important subproblem is learning the knowledge refinement function $R(s, a)$ which captures the OE Tradeoff. Learning this function requires measuring the amount of knowledge refinement produced by various sensing activities dependent on the sensing state, then using these values to generalize a model. The specific measure used to learn this model is dependent on the knowledge framework used by the agent, as well as the domain application. Considering the relationship between this $R(s, a)$ function and resource usage in sensing, we see that $R(s, a)$ is a state-dependent sensing performance profile mapping sensing activities (resource usage) into sensing performance (knowledge refinement). However, such a performance profile is not restricted to be monotonic; thus it can model the Observer Effect, matching the solution requirement set forth in Section 2.

4. CONCLUSION

In conclusion, we have introduced the Observer Effect arising from agent sensing using stateful resources, a novel challenge within bounded rationality. This phenomenon creates a tradeoff between 1) satisfying the need for knowledge refinement, and 2) satisfying the need to avoid knowledge corruption from distorted sensing outcomes intended for knowledge refinement. We model the problem of choosing sensing activities to balance this tradeoff in an active perception setting with the Observer Effect MDP and use RL to learn a controller for choosing sensing activities.

Based on this work, we have identified several important avenues for future work. First, we are currently conducting experiments to explore the OE and evaluate our methodology. Second, we intend to extend our approach to partially observable environments by modeling the decision process instead as a POMDP [4].

5. ACKNOWLEDGMENTS

This research was supported in part by a grant from the Department of Education (grant P200A040150), an NSF grant (DBI0742783), and a NSF Graduate Research Fellowship.

6. REFERENCES

[1] Conlisk, J. 1996. Why bounded rationality? *J. Economic Literature*. 34. June 1996. 669-700.

[2] Fowler, H.J. and Leland, W.E. 1991. Local area network traffic characteristics with implications for broadband network congestion management, *IEEE J. on Selected Areas of Comm.*, 9(7), 1139-1149.

[3] Gigerenzer, G. and Goldstein, D.G. 1996. Reasoning the fast and frugal way: Models of bounded rationality, *Psychological Review*. 103(4). 650-669.

[4] Kaelbling, L.P., Littman, M.L., and Cassandra, A.R. 1998. Planning and acting in partially observable stochastic domains, *AI*, 101, 99-134.

[5] Klein, J., Moon, Y., and Picard, R.W. 2002. This computer responds to user frustration: theory, design, and results. *Interacting with Computers*, 14, 119-140.

[6] Raja, A. and Lesser, V. 2007. A framework for meta-level control in multi-agent systems. *JAAMAS*, *15*, 147–196

[7] Rubinstein, A. 1998. *Modeling Bounded Rationality*. MIT Press: Cambridge, MA.

[8] Sutton, R.S. and Barto, A.G. 1998. Reinforcement learning: an introduction. MIT Press:Cambridge, MA.

[9] Weyns, D., Steegmans, E., and Holvoet, T. 2004. Towards active perception in situated multi-agent systems, *Applied Artificial Intelligence*. 18. 867-883.

[10] Zilberstein, S. 1996. Resource-bounded sensing and planning in autonomous systems. *Autonomous Robots*. 3. 31-48.

[11] Zilberstein, S. 2008. Metareasoning and bounded rationality. *Proc. of AAAI Workshop on Metareasoning: Thinking about Thinking*.